Toward a Calculable Language of Meaning: Integrating Formal Semantics, Cognitive Modeling, and Semantic Entropy

Douglas C. Youvan

doug@youvan.com

October 6, 2024

The concept of an a priori calculable language system represents a significant departure from traditional models of language processing, where meaning is inferred post-hoc through vast amounts of data. This paper explores how a system can be developed to calculate meaning directly from the structure of language, using formal semantic rules, cognitive modeling, and semantic entropy. By embedding meaning within linguistic structures, such a system could revolutionize artificial intelligence, enabling models to reduce their dependence on large datasets and achieve deeper, context-sensitive understanding. Additionally, this approach offers insights into human cognition, shedding light on how the brain dynamically generates meaning through context and sensory experiences. We propose a framework for developing this system, leveraging advances in neural-symbolic integration, dynamic ontologies, and graph neural networks. The result is a language model capable of more efficient communication and richer semantic depth, with profound implications for AI, cognitive science, and future communication technologies.

Keywords: calculable language, formal semantics, cognitive modeling, semantic entropy, neural-symbolic integration, dynamic ontology, graph neural networks, meaning generation, artificial intelligence, embodied cognition, predictive coding, quantum information theory, efficient communication systems, semantic grammar. 51 pages.

I. Introduction

Context and Motivation:

Modern linguistic models, particularly those used in artificial intelligence (AI), rely heavily on statistical inference to understand and generate language. Popular models such as transformers (used in systems like GPT and BERT) extract meaning from massive datasets by identifying patterns in the relationships between words, sentences, and even entire texts. This process of meaning extraction is a posteriori—meaning is inferred after the data has been processed, with the model essentially "guessing" the most likely next word or sentence based on statistical probabilities. While these models have achieved impressive performance in natural language processing (NLP), they rely on an enormous volume of training data to perform well, and the meaning they generate is often based on shallow associations rather than deep, conceptual understanding.

Current AI language models lack the ability to calculate meaning in a systematic, logical manner. They don't "know" the meaning of a sentence or word in the way that a human does; rather, they rely on probability distributions over words or phrases. This leads to limitations such as:

- Contextual ambiguity: Models may struggle with nuanced or ambiguous statements that require deeper conceptual understanding.
- Data dependence: They require vast datasets to learn patterns of language use, and even then, they may not fully grasp complex or less common meanings.
- Surface-level associations: They often rely on superficial associations between words rather than a structured understanding of underlying concepts.

The idea of creating a language where meaning can be calculated a priori—before any external data is encountered—represents a radical departure from this approach. In this vision, meaning is embedded directly into the structural components of the language, allowing it to be derived logically and systematically. Each symbol or combination of symbols would carry intrinsic meaning based on its structural properties, eliminating the need for large datasets or statistical inference to derive meaning.

Such a language could have profound implications for multiple fields:

- Artificial Intelligence (AI): A system that calculates meaning rather than inferring it could lead to more interpretable and reliable AI models, with less reliance on massive data sets and better handling of ambiguity.
- Cognitive Science: A calculable language could offer insights into how humans construct meaning, bridging the gap between symbolic logic and cognitive processes. It could lead to a deeper understanding of how the human brain might "calculate" meaning based on neural structures.
- Computational Linguistics: The development of such a language could push the boundaries of what is possible in natural language processing (NLP), offering new models of communication that are efficient, structured, and semantically rich.

The motivation for this project lies in addressing the limitations of current models and pushing forward a new approach that integrates linguistic structure, cognitive modeling, and information theory. It could unlock new levels of expressiveness, clarity, and efficiency in both human and machine communication.

Objective:

The objective of this paper is to explore the theoretical framework needed to develop a language where meaning is directly embedded into the structure of the language itself. This involves drawing on insights from multiple disciplines—formal semantics, cognitive modeling, and information theory—to build a cohesive model of language where symbols and structures are linked directly to their meanings.

The primary goal is to:

1. Define the structural principles that would allow meaning to be calculated a priori, exploring how formal semantic rules can encode deep, context-rich meanings.
2. Integrate cognitive models to simulate how a language system could reflect human-like processes for constructing meaning, without relying on extensive training data.

3. Develop a new theory of information encoding, where symbols in the language carry not only syntactic but semantic weight, reducing the dependence on large datasets and enabling more efficient, meaningful communication.

By the end of this paper, the intention is to provide a blueprint for future research into the construction of this new language model. Such a system would have far-reaching applications in AI, natural language processing, cognitive science, and even human communication, allowing for deeper, more precise exchange of information in a way that current models cannot achieve.

II. The Limits of Current Language Models

A. Current Approaches to Meaning

1. Modern AI and Statistical Meaning Inference

Most contemporary AI models, particularly those used in natural language processing (NLP), infer meaning through statistical associations between words, phrases, and sentences in vast datasets. The dominant models in this field are transformers (e.g., GPT, BERT), which rely on a neural network architecture that tracks dependencies between tokens (words or subwords) in a sequence. These models work by predicting the next word in a sentence or completing a sentence based on patterns learned from massive training corpora.

For example, transformers achieve state-of-the-art performance in various NLP tasks by employing mechanisms like self-attention, which allows the model to weigh the importance of different words in a sentence based on their relationships to each other. However, the underlying process remains probabilistic: the model has no intrinsic understanding of the meaning of a sentence. Instead, it calculates the likelihood of different word sequences based on statistical regularities found in the data.

This statistical approach to meaning has several key limitations:

- Surface-Level Understanding: These models often capture the superficial associations between words but miss deeper, abstract, or metaphorical

meanings. This leads to issues when dealing with ambiguous or complex texts where the context isn't straightforward.

- Data Dependency: AI models rely on extremely large corpora to learn the patterns of language use, which means their ability to understand or generate coherent meaning is constrained by the quality, size, and diversity of the data. Without large-scale data, models perform poorly and struggle to generalize to new contexts.
- Contextual Fragility: While transformers handle local context well, they sometimes falter when dealing with long-term dependencies or global context, as the associations between words become less direct the further apart they are in a sentence or paragraph.

In essence, while modern AI can generate impressive text outputs, its "understanding" of language remains fundamentally statistical and lacks a robust, internally calculated notion of meaning. It cannot derive meaning without pre-existing data, which limits its ability to handle truly novel, ambiguous, or complex situations where deep understanding is required.

2. Formal Semantics: Defining Meaning Post-Hoc

In contrast to the statistical models used in AI, formal semantics aims to define meaning through rigorous logical structures. Formal semantics uses tools like truth conditions to define the meaning of a sentence based on whether it accurately reflects a state of affairs in the world. For instance, the meaning of a sentence like "The cat is on the mat" is determined by whether there is, in fact, a cat on the mat.

Formal semantics provides a post-hoc method of deriving meaning by mapping linguistic expressions to entities in the real world or in abstract logical structures. Systems like model-theoretic semantics and truth-conditional semantics use these mappings to describe the meaning of sentences in terms of possible worlds or logical models. While this approach is rigorous and provides a solid theoretical foundation for understanding meaning, it has its own set of limitations when applied to real-world language processing:

- Rigid Structure: Formal semantics relies on well-defined models of the world, which makes it difficult to apply to more fluid or ambiguous natural

language scenarios. Everyday language often involves vagueness, metaphors, or incomplete information, which formal systems struggle to account for.

- Non-Computational Focus: Although formal semantics provides deep insights into how meaning can be logically structured, it is not designed to be computationally efficient or scalable for real-time language processing, making it less applicable for AI systems that need to generate or interpret meaning dynamically.

Thus, formal semantics, while theoretically powerful, is often impractical for large-scale language processing and is rarely used in its pure form in modern AI systems. Instead, meaning remains something that is inferred through statistical patterns, not calculated directly through logical structures.

B. The Bandwidth-Meaning Paradox

1. Shannon's Information Theory and Bandwidth

Claude Shannon’s information theory, which is the bedrock of modern communication systems, focuses on the efficient transmission of data rather than the meaning of that data. According to Shannon, the amount of information transmitted is measured by its entropy—a concept that describes the unpredictability or uncertainty of the transmitted message. The more symbols (or characters) a communication system uses, the greater its bandwidth or its ability to transmit information efficiently.

For example, a system that uses a binary alphabet (0 and 1) has a certain capacity to encode and transmit information. A system that uses a larger alphabet, such as one with 26 characters (like the English alphabet), can encode more information per symbol, allowing for a greater transmission rate (i.e., increased bandwidth).

However, Shannon’s theory is entirely agnostic to meaning. The focus is on how much information is being transmitted, not what that information means. This leads to what we might call the bandwidth-meaning paradox:

- Increased Bandwidth ≠ Increased Meaning: As we increase the size of the alphabet or the complexity of the encoding system, we increase the

amount of information that can be transmitted. But this does not necessarily correlate with the depth, richness, or complexity of the meaning that is conveyed. For example, encoding a random string of characters with a large alphabet increases entropy, but those characters may not carry any meaningful information.

Shannon's model views language as a sequence of independent symbols, focusing purely on signal transmission. Meaning, in this framework, is irrelevant to the transmission process, which is problematic when applied to systems where meaning is the key focus, such as human communication or AI understanding.

2. Larger Alphabets and Semantic Complexity

While larger alphabets can increase symbolic complexity and transmission efficiency, they don't automatically translate into greater semantic complexity. For instance, Chinese logograms may represent more information per character than English letters, but this does not mean that a system using logograms is intrinsically more "sophisticated" than one using an alphabet.

This raises the question of how we can increase semantic density—the amount of meaning conveyed per unit of information—rather than just bandwidth. Shannon's theory doesn't offer tools for calculating the semantic content of messages; it only quantifies the efficiency of transmission. Therefore, if our goal is to create a system that conveys meaning more effectively, we must look beyond traditional information theory.

In a language system that calculates meaning a priori, the symbols themselves would be designed to carry semantic weight, not just transmission efficiency. This would involve integrating the semantic properties of symbols with the rules of the language to ensure that meaning is both encoded and transmitted. Such a system would bypass the bandwidth-meaning paradox by making meaning an intrinsic part of the encoding process, not an emergent property of statistical patterns or transmission efficiency.

III. Formal Semantics and Compositional Structures

A. The Role of Formal Semantics

Formal semantics aims to rigorously define how language relates to meaning using mathematical and logical systems. The primary goal is to specify how sentences and phrases in natural language correspond to states of the world, abstract concepts, or logical relationships. Formal semantic theories like truth-conditional semantics, model-theoretic semantics, and possible-worlds semantics have provided strong frameworks for understanding meaning in a structured way.

1. Existing Formal Semantic Systems

- Truth-Conditional Semantics: This theory defines the meaning of a sentence by specifying the conditions under which it would be true. For example, the sentence "The sky is blue" is true if and only if the sky is blue in the world we're referencing. Truth-conditional semantics works well for simple, factual statements but struggles with more complex or metaphorical meanings, such as those found in figurative language, indirect speech, or emotional expressions.
- Model-Theoretic Semantics: Here, sentences are evaluated within formal models of the world, consisting of entities, relationships, and functions. Each linguistic expression is mapped onto objects or relationships within this model, making it easier to systematically reason about meaning. However, the limitations of model-theoretic semantics lie in its dependence on pre-specified models of the world, which do not always account for the fluidity or ambiguity of real-world language use.
- Possible-Worlds Semantics: In this approach, meaning is evaluated across a range of possible worlds or hypothetical scenarios. This allows for handling modal expressions like "could," "should," or "might," which rely on considering alternative realities. While powerful for evaluating counterfactuals or hypothetical reasoning, possible-worlds semantics also struggles with real-time language processing and can be cumbersome for AI systems to implement efficiently.

2. Limitations in A Priori Meaning Calculation

Despite their rigor, existing formal semantic systems are fundamentally a posteriori. They map linguistic expressions to pre-defined logical structures or world models, requiring interpretation and evaluation after the fact. This approach limits their ability to calculate meaning directly from the structure of language itself. The meaning must be derived based on external reference points, such as real-world facts or logical models, rather than being intrinsic to the linguistic components.

For example, in truth-conditional semantics, the meaning of a sentence is dependent on external verification. In model-theoretic semantics, the meaning is determined by the relationships within a predefined model. Neither of these approaches can calculate meaning independently of an external system for evaluation. This reliance on external factors poses a challenge for any system that aims to embed meaning directly into the language structure.

To move toward a language where meaning is intrinsic and calculable a priori, a new framework is required—one that integrates meaning into the relations between linguistic elements themselves, rather than treating meaning as something that must be evaluated after language is constructed. This is where compositionality and category theory come into play.

B. Compositionality and Category Theory

1. The Principle of Compositionality

Compositionality is a key concept in semantics and linguistics. It states that the meaning of a complex expression (such as a sentence) is a function of the meanings of its parts and the way those parts are syntactically combined. For example, the meaning of "The cat is on the mat" is determined by the meanings of "cat," "mat," "is on," and their arrangement.

In traditional formal semantics, compositionality is used to ensure that meaning can be built up incrementally from smaller components, such as words or phrases, into larger structures like sentences or paragraphs. However, traditional

systems of compositionality still rely on external models or logical conditions to calculate meaning post-hoc.

The challenge is to extend compositionality in such a way that meaning emerges intrinsically from the structure of the language components, without needing external reference points. This is where category theory, a branch of mathematics that deals with abstract structures and relationships, could offer a new approach.

2. Category Theory: A New Framework for Meaning

Category theory provides a formal system where relationships between objects are as important as the objects themselves. In category theory, we focus on morphisms (the relationships between objects) rather than just the objects alone. This shift from "things" to "relationships" could be the key to embedding meaning directly into the structural relationships between linguistic components.

In a language system based on category theory:

- Objects can represent words, phrases, or even entire sentences.
- Morphisms represent the syntactic or logical relationships between these objects.

Category theory can provide a compositional framework where the meaning of complex expressions is not only a function of the meanings of their parts but also emerges from the morphisms that connect them. This allows for meaning to be calculated based on the relationships between components, without needing external verification.

3. Mapping Sentence Structures to Logical Relationships

In traditional language models, sentence structures are often represented in terms of syntax trees or dependency graphs. These structures describe how words relate to one another grammatically, but they don’t typically include semantic relationships. Using category theory, we can define a system where sentence structures are mapped directly to logical or ontological relationships. Here’s how it could work:

- Example: Simple Sentence in Category Theory
    - Consider the sentence "The cat chases the mouse."

- In a traditional compositional system, the meaning would be derived from the individual meanings of "cat," "chases," and "mouse," plus the grammatical structure of the sentence.
- In a category-theoretic approach, we represent the entities "cat" and "mouse" as objects, and "chases" as a morphism—a relationship that connects the two objects. The meaning of the sentence emerges from this morphism, as it defines a specific type of relationship (chasing) between the two objects.
- The structure and the morphism together encode not just the fact that "chasing" connects "cat" and "mouse," but also additional semantic layers: for instance, whether the action is ongoing (present tense), whether it's hypothetical (modal), or whether it refers to a specific instance (definiteness).

By treating meaning as something that emerges from these morphisms—not just from the individual components—the system becomes more calculated and intrinsic. This means that meaning can be calculated a priori, as the relationships between components carry semantic weight directly within the system, without the need for external reference points or post-hoc evaluations.

4. Category Theory as a Foundation for an A Priori Language System

A category-theoretic approach to compositionality would allow us to design a language where meaning is encoded within the structural relationships between components. In such a system:

- Words and sentences are objects.
- Grammatical relationships are morphisms.
- Meaning emerges as a natural consequence of how objects and morphisms are composed.

This would offer a powerful alternative to current formal semantic models by allowing for meaning to be derived intrinsically through the structure of the language itself. Such a system could be more efficient, flexible, and precise, enabling the calculation of meaning in real-time without requiring external data or models.

IV. Cognitive Modeling: Mimicking Human Meaning Construction

A. Cognitive Modeling of Language Understanding

Cognitive modeling aims to simulate how humans process and understand language by creating computational or theoretical models that approximate human cognition. These models attempt to replicate the processes by which the brain constructs meaning from sensory inputs, words, sentences, and larger discourse. Understanding the human cognitive mechanisms involved in language comprehension provides valuable insights into how we might design artificial systems that construct meaning in a similar way.

1. The Role of Cognitive Models in Meaning-Making

Human language comprehension is a complex, dynamic process that involves multiple cognitive systems working together:

- Perceptual processing: Our brains use sensory data (auditory, visual, etc.) to interpret the raw forms of linguistic input, such as spoken words or written symbols.
- Lexical retrieval: Upon hearing or seeing a word, the brain retrieves its associated meaning from a vast mental lexicon, mapping the sound or symbol to its semantic content.
- Contextual integration: As each word or phrase is understood, it is integrated into the surrounding linguistic context, taking into account the meaning of previous words and the larger discourse.
- Conceptual reasoning: To fully grasp the meaning of a sentence or conversation, humans rely on their prior knowledge of the world, as well as inferential reasoning, to fill in gaps and resolve ambiguities.

Several cognitive models have been proposed to explain how the brain processes language:

- Connectionist Models: These models are inspired by neural networks and emphasize the parallel, distributed nature of language processing. They explain language understanding as the result of pattern recognition across

vast networks of interconnected neurons, with meaning emerging from the statistical relationships between different linguistic units.

- Symbolic Models: These models propose that the brain processes language using a set of rules, much like a computer program. In symbolic models, meaning is encoded in discrete symbols (words, phrases) and processed according to syntactic and logical rules.
- Hybrid Models: A more recent approach combines elements of both connectionist and symbolic models, recognizing that meaning arises from both the statistical regularities in language (connectionist) and the structured, rule-based nature of linguistic relationships (symbolic).

These cognitive models underscore that human understanding of meaning is an active, interpretative process. We don't simply decode static information from language; we constantly adjust our understanding based on new inputs, prior knowledge, and contextual clues. This adaptive process is one of the key differences between human cognition and current AI models, which rely more on statistical pattern-matching than dynamic, real-time meaning construction.

2. Simulating Human Language Understanding in AI

In the context of AI and natural language processing, most current systems attempt to simulate the human process of meaning-making, but they do so in a simplified or limited way. AI systems use techniques like:

- Word embeddings: Words are represented as vectors in a multi-dimensional space, where proximity to other vectors represents semantic similarity. While this allows AI to simulate some aspects of human understanding, it falls short of true meaning construction, as the relationships between words are based solely on co-occurrence patterns, not on deeper conceptual or experiential meaning.
- Sequence models: Neural networks, such as recurrent neural networks (RNNs) and transformers, model language as a sequence of tokens, attempting to predict the next token or generate a coherent sequence based on prior tokens. While these models can produce coherent output, they don't "understand" meaning in the way humans do—they simply reproduce statistical patterns from training data.

To build a system where meaning is calculated a priori, we would need to go beyond statistical inference and simulate the dynamic, context-sensitive meaning construction process seen in human cognition. This is where predictive coding comes in as a promising framework for designing such a system.

B. Predictive Coding and Meaning

Predictive coding is a leading theory in neuroscience that suggests the brain operates as a prediction machine. According to this theory, the brain continually generates predictions about incoming sensory data, and these predictions are updated in response to actual sensory input. In this framework, the brain is constantly trying to minimize the discrepancy between its internal model of the world and the external reality it perceives, a process known as prediction error minimization.

1. The Predictive Coding Model

In predictive coding, the brain works as a hierarchical system, where higher levels generate predictions about what is expected at lower levels, and sensory input is compared to these predictions. The prediction error—the difference between the expected and actual sensory input—is then fed back up the hierarchy to adjust the internal model. This feedback loop allows the brain to continuously refine its understanding of the world, resulting in more accurate predictions over time.

- Top-Down Predictions: Higher cognitive layers (involving abstract reasoning, memory, and prior knowledge) send predictions to lower sensory layers about what inputs to expect (e.g., anticipating the next word in a sentence based on context).
- Bottom-Up Sensory Input: Lower layers provide real-time sensory input (e.g., the actual words being heard or read), which are compared against the top-down predictions.
- Error Correction: If the sensory input doesn’t match the prediction, the brain adjusts its internal model to reduce the prediction error, effectively "learning" from the discrepancy.

Predictive coding thus provides a powerful model for understanding how humans process language and construct meaning dynamically. Rather than passively receiving and decoding language, the brain actively predicts and updates its understanding in real-time.

2. Applying Predictive Coding to a Language System

If we want to design a language system that calculates meaning a priori, we can leverage predictive coding principles to create a dynamic, self-updating model of meaning. In such a system, symbols (words, phrases) would generate meaning through continuous internal updates, just as the brain adjusts its internal model based on new inputs.

Here's how a predictive coding-based language system could work:

- Initial Prediction: When a word or sentence is introduced into the system, an initial prediction about its meaning is generated based on the system's existing knowledge or context. This prediction is not final but serves as a starting point.
- Feedback Loop: As more symbols (words) are processed, the system compares the meaning generated by each new input to its prior predictions. If the new input contradicts the predicted meaning, the system adjusts its internal representation to reduce the prediction error. For example, if the sentence structure changes or an unexpected word is introduced, the system updates its understanding to reflect the new context.
- Contextual Integration: The system doesn't treat each symbol or word as an isolated unit. Instead, it integrates new information into a broader contextual framework, where the meaning of each word is influenced by the surrounding words and the overall discourse. This mimics how humans rely on context to interpret ambiguous or incomplete sentences.
- Meaning as a Dynamic Process: In this model, meaning is not static or pre-defined but is constantly evolving based on new inputs and the system's prior knowledge. Just as the human brain refines its understanding over time, the system would update its internal representation of meaning as it encounters new linguistic data.

3. Example: Predictive Coding in Sentence Processing

Consider the sentence: "The cat chased the..." A predictive coding-based system would initially generate a set of predictions for the next word (e.g., "mouse," "dog," etc.) based on the context provided by the previous words. When the actual next word ("mouse") is processed, the system compares this input to its prediction. If the prediction was correct, the system reinforces its internal model. If the prediction was incorrect, the system adjusts its model to accommodate the new input.

Now, consider a more complex scenario: "The cat chased the mouse, but suddenly..." The addition of "but suddenly" introduces a change in the expected outcome, forcing the system to revise its internal prediction about what comes next. This continuous process of prediction and error correction allows the system to handle complex, dynamic language inputs more effectively than traditional models.

4. Advantages of a Predictive Coding-Based Language System

- Dynamic, Real-Time Meaning Construction: Just as the brain constructs meaning dynamically in response to new inputs, a predictive coding system could continuously update its meaning representations, allowing for real-time, context-sensitive language processing.
- Handling Ambiguity and Complexity: By incorporating feedback loops and error correction, the system would be better equipped to handle ambiguous or complex sentences, adjusting its understanding as new information becomes available.
- Reduced Data Dependency: Instead of relying on vast datasets to infer meaning from statistical patterns, a predictive coding system could construct meaning from a smaller set of well-defined rules and prior knowledge, making it more efficient and less data-hungry.

V. A New Theory of Semantic Information Encoding

A. Moving Beyond Shannon's Entropy

1. Shannon's Entropy and Its Limitations

Claude Shannon's information theory, a groundbreaking framework in communication and signal processing, fundamentally transformed how we understand the transmission of information. Shannon's **entropy** measures the amount of uncertainty or unpredictability in a set of possible messages. Mathematically, entropy $H(X)$ for a discrete random variable X with probabilities $p(x_i)$ is given by:

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_b p(x_i)$$

This formula quantifies the average amount of information (or surprise) produced by a source. In simple terms, Shannon's entropy tells us how efficiently a message can be compressed or transmitted over a communication channel, based on the likelihood of different symbols appearing.

However, Shannon's entropy doesn't account for meaning—it focuses purely on the unpredictability of symbols without any regard for the semantic content of the messages. A highly unpredictable message (one with high entropy) could consist of random symbols that carry little or no actual meaning, while a message with low entropy could be rich with meaningful content. For example, a long string of random letters may have high entropy but convey no useful information, whereas a concise sentence like "Water is essential for life" has low entropy but is semantically rich.

Thus, while Shannon's model is useful for understanding the efficiency of symbol transmission, it is agnostic to the depth, complexity, or richness of the meaning conveyed. To address this, we need to move beyond Shannon's entropy and introduce a concept of semantic entropy that accounts for the complexity and richness of meaning encoded in a message.

2. Semantic Entropy: A New Metric for Meaning

Semantic entropy can be thought of as a measure of the meaningful complexity of a message, rather than just its unpredictability. The goal of semantic entropy is to

quantify how much meaningful information is packed into a message, taking into account not only the arrangement of symbols but the relationships between them, the context they create, and the depth of understanding they evoke.

To formalize this, we might introduce a new equation or framework for measuring semantic complexity. Consider the following components:

- Contextual Information: Unlike Shannon's entropy, which treats each symbol as independent, semantic entropy must account for contextual dependencies between symbols. The meaning of a word or sentence often depends on its relationship to surrounding words, sentences, or even larger discourse structures.
- Symbol-Concept Mapping: In semantic entropy, symbols (words, phrases, etc.) would map onto concepts or ideas rather than just being raw symbols. The richness of meaning could be determined by the complexity of the conceptual relationships between these symbols.
- Hierarchical Structure: Semantic entropy should account for the hierarchical nature of meaning in language. Sentences are built from words, paragraphs from sentences, and so on, with each level contributing to the overall meaning.

A possible formula for semantic entropy might look something like this:

A possible formula for **semantic entropy** might look something like this:

$$H_s(M) = -\sum_{i=1}^{n} p(c_i) \log_b p(c_i) \cdot W(c_i)$$

Where:

- $H_s(M)$ is the semantic entropy of a message M,
- $p(c_i)$ is the probability of the occurrence of a concept c_i,
- $W(c_i)$ is a weight that reflects the **semantic richness** or **complexity** of the concept c_i.

In this model, concepts are weighted not just by their frequency but by their depth—more abstract or nuanced concepts would contribute more to the overall

semantic entropy than simpler, more concrete concepts. This allows the system to differentiate between shallow messages and deeply meaningful ones.

3. Measuring Semantic Complexity

One of the key challenges of this new framework is how to objectively measure the richness or complexity of meaning. Some possible approaches could include:

- Ontological Depth: A concept's position within an ontological hierarchy (e.g., more abstract concepts like "freedom" or "justice" would have higher weights than concrete concepts like "cat" or "table").
- Contextual Variability: Concepts that shift meaning significantly based on context (e.g., metaphors, idiomatic expressions) could be given higher weights, as they require more cognitive processing to interpret.
- Interconnectedness: Concepts that are highly interconnected with other ideas or concepts (e.g., those that appear frequently in many contexts) could also carry more weight, as they contribute to a more nuanced understanding of the overall message.

By developing a model of semantic entropy that incorporates these factors, we can begin to quantify not just how much information a message contains, but how meaningfully complex that information is. This would represent a significant step forward from Shannon's model, allowing us to encode not only the efficiency of transmission but the depth of the meaning being communicated.

B. Quantum Information as a Model

While Shannon's information theory is focused on the classical transmission of symbols, quantum information theory deals with the transmission and processing of information in quantum systems. Quantum systems are unique in that they allow for phenomena like superposition (where a system exists in multiple states simultaneously) and entanglement (where the state of one particle is dependent on the state of another, even at a distance).

These properties of quantum systems offer a potential model for encoding layered and context-rich meanings in language. By leveraging ideas from quantum

information theory, we can explore how semantic meaning might be encoded in ways that reflect the complexity and interdependence of concepts.

1. Superposition and Layered Meaning

In quantum systems, a particle can exist in multiple states at once (superposition) until it is observed or measured. This property of quantum systems has an intriguing parallel in language, where words or sentences can carry multiple layers of meaning simultaneously. For example, a word like "light" can refer to both the physical phenomenon of illumination and a metaphorical sense of enlightenment or understanding.

A quantum-inspired model for language could encode superpositions of meaning, where each word or sentence carries multiple potential meanings that depend on the context in which it is used. The process of interpreting the word would then be analogous to the collapse of a quantum superposition, where one specific meaning is selected based on the surrounding context.

In this framework:

- Words or symbols would be treated as quantum states with the potential to encode multiple meanings simultaneously.
- The act of interpreting a message would involve "collapsing" the superposition into a single, contextually appropriate meaning, based on the relationships between the symbols and the broader context in which they appear.

This quantum-inspired approach could allow for more efficient encoding of ambiguous or layered meanings, where a single word or phrase can carry multiple interpretations without needing to explicitly specify each one.

2. Entanglement and Contextual Dependence

In quantum systems, entanglement occurs when two or more particles become linked in such a way that the state of one particle is dependent on the state of the other, even across large distances. This concept has an interesting parallel in language, where the meaning of one word or phrase often depends on the meaning of another word or phrase.

For example, consider the sentence "The cat is out of the bag." The phrase "out of the bag" only makes sense when understood in the context of the entire sentence and the idiomatic meaning of the phrase as a whole. This contextual dependence is similar to quantum entanglement, where the state of one particle cannot be understood in isolation but must be considered in relation to its entangled counterpart.

In a quantum-inspired semantic system:

- Words or phrases would be treated as entangled states, where their meanings are dependent on their relationships to other words or phrases in the message.
- The meaning of a word would be determined not just by its individual properties but by its entangled relationships with other words in the sentence or discourse.

This model would allow for a more holistic understanding of meaning, where the relationships between words carry as much weight as the words themselves. Just as entangled particles can exhibit complex behaviors that are not reducible to their individual states, entangled meanings in language would reflect the intricate web of relationships between concepts.

3. Quantum States for Deep Context Encoding

Quantum information theory also introduces the idea of quantum states, which can carry much more information than classical bits. A classical bit can exist in one of two states (0 or 1), while a quantum bit (or qubit) can exist in a superposition of both 0 and 1, allowing for the encoding of more complex information.

In a language system, quantum states could be used to encode deep context or multi-dimensional meanings. Instead of representing each word or concept as a discrete, static entity, we could represent it as a multi-dimensional quantum state that evolves and changes depending on the context. This would allow for the encoding of richer, more nuanced meanings that reflect the full complexity of human communication.

VI. Deep Integration of Meaning and Structure

A. Semantic Grammar

1. The Concept of Semantic Grammar

In traditional grammar, the rules of syntax and morphology are primarily concerned with the structure of language—how words are arranged into sentences and how word forms change depending on tense, number, gender, and so on. The primary focus is on the form of language rather than its meaning. However, the idea of semantic grammar shifts the focus from mere structure to a system where each grammatical element also carries semantic weight—meaning becomes inseparable from structure.

In a semantic grammar, the rules of syntax, word order, and morphology would not only dictate how sentences are formed but also how meaning is constructed in real-time. In this model:

- Syntax (the arrangement of words) directly encodes semantic relationships, such as subject-object relationships, causality, or temporal sequences.
- Word order is not just a syntactic rule but also reflects semantic priority or focus. For instance, shifting word order could convey changes in meaning, such as emphasis or thematic importance, rather than merely grammatical correctness.
- Morphology (the inflections and modifications of words) would encode semantic nuances such as aspect, intention, or modality directly into the language. For example, verb tenses would not just denote time but could indicate the certainty or hypothetical nature of an event.

This approach mirrors functional grammar theories, which argue that the function of language is primarily communicative and that linguistic structure reflects how humans process and express meaning. Semantic grammar would go a step further by ensuring that meaning is an intrinsic part of every grammatical rule, allowing meaning to be calculated as the sentence is constructed, rather than inferred from the completed structure.

2. Neural-Symbolic Systems for Meaning Integration

To implement a semantic grammar in AI systems, we could turn to neural-symbolic integration, a field that combines the strengths of symbolic logic (with its rule-based precision) and neural networks (with their ability to handle large-scale, dynamic patterns).

- Symbolic logic excels in explicitly encoding knowledge and rules, such as the syntactic and semantic rules that could govern a semantic grammar. These systems provide clear, interpretable representations of how language should work, making them ideal for capturing the rigid structure needed for a grammar system.
- Neural networks, on the other hand, excel in pattern recognition and can handle the variability, flexibility, and ambiguity present in natural language. They can learn from data and adapt to different linguistic inputs, which is crucial for dealing with the complexities of real-world language.

By combining these approaches, we can create a neural-symbolic system where meaning is embedded directly into the linguistic structures. The symbolic component would provide the grammatical rules and the semantic relationships, while the neural component would ensure flexibility and adaptability, learning how these rules apply in different contexts.

For example, a neural-symbolic system implementing semantic grammar could work as follows:

- Word-level encoding: Each word or phrase is represented as a symbolic token that carries both its syntactic category (e.g., noun, verb) and its semantic meaning (e.g., agent, action, object). The neural network would process these tokens, recognizing patterns and relationships that map these symbols to their meaning.
- Sentence-level composition: As words are combined into sentences, the symbolic rules of semantic grammar ensure that the structure of the sentence directly reflects the relationships between its components (e.g., who is doing what to whom). The neural network continuously refines its understanding, adjusting the interpretation of meaning based on both the rules and the context.

This approach allows the calculation of meaning to occur as the language is processed—meaning is not inferred after the fact but rather emerges naturally from the structure of the language as it is constructed. In this way, neural-symbolic systems could integrate meaning into structure, making AI systems more precise and human-like in their interpretation of language.

3. Example: Calculating Meaning in a Semantic Grammar

Consider the sentence: "The dog chased the cat."

In a traditional system, the meaning of this sentence would be inferred based on the word meanings (dog, chase, cat) and their syntactic relationships (subject-verb-object). In a semantic grammar system, the structure of the sentence itself would carry semantic weight:

- The subject-verb-object structure tells us that the dog (the agent) is performing the action (chasing) on the cat (the patient).
- The word order tells us that the dog is the active entity, and the cat is the one being affected by the action.
- If we shift the word order ("The cat was chased by the dog"), the grammatical structure changes, but so does the semantic focus—the cat becomes the focus, even though the event remains the same.

The neural-symbolic system would encode this information dynamically, adjusting its internal representation of meaning as the sentence unfolds. As a result, meaning is constructed in real-time, with each word and grammatical structure contributing directly to the system’s understanding.

B. Embodied Cognition and Conceptual Frames

1. The Embodied Cognition Model

Embodied cognition is a theory in cognitive science that argues that cognitive processes, including language comprehension, are deeply rooted in the body's interactions with the physical world. In this view, meaning is not an abstract, disembodied concept but is instead grounded in sensory and motor experiences. For example:

- The word "grasp" is understood not only in its abstract sense (to grasp an idea) but also through our physical experience of using our hands to grasp objects.
- The word "run" invokes a mental simulation of the act of running, drawing on our bodily experience of movement and exertion.

This theory challenges the traditional view of language as purely symbolic and suggests that meaning is closely tied to our physical experiences in the world. Embodied meaning emerges through conceptual metaphors and frames, which link abstract concepts to concrete, bodily experiences.

2. Conceptual Metaphors and Frames

Conceptual metaphors are a key mechanism through which embodied cognition operates. According to George Lakoff and Mark Johnson, we often use metaphorical structures to understand abstract concepts by linking them to more concrete, bodily experiences. For example:

- We talk about time as if it were space: "We are approaching the deadline" or "We're running out of time."
- We conceptualize understanding as grasping: "I can't grasp that idea."

Frames are cognitive structures that shape how we interpret new information based on past experiences. For instance, when we hear the word "hospital," we activate a mental frame that includes doctors, nurses, patients, and medical equipment, all of which help us understand related language in context.

These metaphors and frames provide contextual scaffolding for meaning, allowing us to understand abstract ideas through the lens of concrete experiences. In a language system that incorporates embodied cognition, the meanings of words and sentences would be calculated based on the conceptual metaphors and frames that link language to sensory and motor experiences.

3. Building a Language System with Embodied Cognition

To develop a language system that calculates meaning based on embodied cognition, we could design the system to:

- Link words and concepts to physical experiences: The system would be trained to associate specific words or phrases with sensorimotor experiences. For example, words like "push" or "pull" would activate representations of physical actions, while more abstract words like "decide" could be linked to experiences of deliberation or choice.
- Use conceptual metaphors: The system would employ metaphorical reasoning to calculate meaning. For instance, when processing a phrase like "running out of time," the system would activate the conceptual metaphor of time as a finite resource, linking the phrase to experiences of scarcity or urgency.
- Activate frames: When encountering certain words or phrases, the system would activate frames based on prior knowledge. For example, the word "journey" could activate a frame that includes notions of movement, obstacles, and destinations, helping the system calculate meaning within this broader context.

4. Example: Calculating Meaning Through Embodied Cognition

Consider the sentence: "He grasped the idea quickly."

In an embodied cognition-based system:

- The word "grasped" would not only be understood in its metaphorical sense (understanding an idea) but would also be linked to the physical experience of grasping something with the hands.
- The system would use the conceptual metaphor "understanding is grasping" to bridge the abstract concept of understanding with the physical experience of grasping.
- The word "quickly" would evoke bodily experiences of speed and movement, adding nuance to the understanding of how the idea was comprehended.

The result is a more holistic calculation of meaning, where the system integrates both the literal and metaphorical aspects of the sentence, drawing on embodied experiences to enhance its understanding.

VII. Ontologies and Self-Organizing Meaning Systems

A. Dynamic Ontologies for Meaning Calculation

1. Traditional Ontologies in Language and Knowledge Representation

In both natural language processing (NLP) and knowledge representation, an ontology is a structured framework that defines the relationships between concepts within a specific domain. Ontologies typically consist of entities (such as objects, ideas, or processes) and the relationships between them, often arranged hierarchically. For example, in a medical ontology, terms like "doctor," "patient," and "treatment" would be connected by relationships that specify how they interact (e.g., a doctor provides treatment to a patient).

Traditional ontologies play an important role in semantic web technologies, knowledge graphs, and other systems where understanding and representing complex relationships between concepts is crucial. They provide:

- Formal structures for defining concepts and their interrelations.
- A way to enable semantic reasoning, where systems can infer new information based on the defined relationships.
- A basis for querying knowledge in a consistent and logical way.

However, traditional ontologies are often static—meaning they are manually created and updated by human experts. This makes them inflexible and slow to adapt to new knowledge or changes in understanding. Furthermore, they rely heavily on human-defined hierarchies and relationships, which may not always reflect the fluid and dynamic nature of meaning in natural language or real-world contexts.

2. The Need for Dynamic, Self-Updating Ontologies

A dynamic ontology is one that can adapt, evolve, and update itself in response to new information. Instead of relying on pre-defined, rigid relationships between concepts, a dynamic ontology would be able to:

- Learn and adapt to new data or contexts in real-time, modifying its internal structure to accommodate changes.

- Automatically integrate new concepts or terms as they emerge, without requiring manual intervention.
- Reorganize relationships based on updated information, making the system more responsive to shifts in meaning or new discoveries.

Such a system would be crucial for real-time meaning calculation because it allows the ontology to continuously adjust as new inputs are processed, ensuring that the system remains flexible and context-sensitive. Unlike traditional ontologies that require human updates, a dynamic ontology would be able to self-organize, creating and refining its own structure based on the data it encounters.

For example, imagine a language system processing news articles in real-time:

- A traditional ontology might struggle to incorporate new entities (such as a newly coined term or emerging concept) without manual updates.
- A dynamic ontology, on the other hand, would recognize patterns in the data and automatically infer new relationships between concepts, allowing it to incorporate new terms or ideas without needing explicit updates from a human curator.

3. Ontologies and Real-Time Meaning Calculation

To calculate meaning in real-time, a dynamic ontology would continuously adjust its structure based on the relationships between concepts as they are encountered in language. This could involve:

- Re-weighting relationships: As new data is processed, the system could adjust the strength of relationships between concepts, placing more emphasis on relationships that are frequently encountered or contextually important and downplaying those that are less relevant.
- Introducing new nodes and edges: As the system encounters new terms or concepts, it could automatically create new nodes (representing concepts) and edges (representing relationships) in the ontology. This allows the system to expand its knowledge base dynamically, without needing human intervention.

- Context-sensitive adjustments: The system could modify its relationships based on the context in which concepts are used. For example, the relationship between "virus" and "vaccine" might be different in the context of a medical discussion than in a biological research context, and the ontology would adjust accordingly.

These dynamic updates allow the system to remain responsive to new information, ensuring that meaning is calculated based on the most up-to-date and relevant relationships between concepts. The system could process language in a way that reflects the fluidity of human meaning-making, where concepts and their relationships shift based on context, culture, and evolving knowledge.

B. Graph Neural Networks

1. The Role of Graph Neural Networks in Structuring Knowledge

Graph Neural Networks (GNNs) represent a powerful tool for encoding and processing relational data. In a GNN, nodes represent entities (such as words, concepts, or objects), and edges represent the relationships between these entities. The structure of the graph captures how knowledge is interconnected, with each node and edge contributing to the overall meaning encoded in the network.

Unlike traditional neural networks that process inputs as vectors, GNNs are designed to operate on data that is represented in graph form, which makes them particularly well-suited for knowledge representation and semantic reasoning. GNNs are capable of:

- Learning relational structures: By processing the connections between nodes, GNNs can learn complex relationships and dependencies between concepts.
- Propagating information: In a GNN, information can propagate across the network, allowing meaning to emerge from network-wide interactions rather than from isolated entities.
- Handling dynamic data: GNNs are inherently flexible, allowing them to adjust as new nodes or edges are introduced. This makes them ideal for

self-organizing systems, where the structure of knowledge must evolve over time.

2. How GNNs Enable Emergent Meaning

The core strength of GNNs is their ability to encode relational information directly into the network's structure. In a language system, each word, phrase, or concept would be represented as a node, and the relationships between these concepts (such as semantic associations, syntactic dependencies, or contextual co-occurrences) would form the edges. The GNN processes the graph in a way that allows meaning to emerge from the interactions between nodes.

- Edge-based learning: GNNs allow each node to learn from the information encoded in its neighboring nodes. For example, in a sentence like "The doctor treats the patient," the node representing "doctor" would learn from its neighboring nodes ("treats" and "patient"), allowing it to update its internal representation based on the relationships it shares with those concepts.
- Contextual propagation: As information propagates through the network, meaning is not just confined to individual words or phrases but is shaped by the contextual dependencies between them. The system would account for the entire graph structure when calculating meaning, ensuring that context-sensitive meanings emerge naturally.
- Node and edge updates: GNNs are dynamic, meaning that as new data is introduced, the network can automatically adjust by adding new nodes or edges, or by re-weighting the importance of existing connections. This allows the system to self-organize as it encounters new language inputs, keeping the graph structure up-to-date with the latest information.

In this way, meaning becomes an emergent property of the network, rather than something that is statically encoded into individual symbols. The structure of the graph and the interactions between its components drive the system's ability to generate context-aware and dynamic interpretations of meaning.

3. Example: Calculating Meaning with a GNN

Consider a sentence like: "The scientist discovered a new particle at the collider." In a GNN-based system:

- Each word or concept (scientist, discover, particle, collider) is represented as a node.
- The relationships between these words (e.g., the action of discovery, the association with the collider, etc.) are encoded as edges.
- As information propagates through the graph, the system learns the meaning of each node based on its connections to other nodes. For example, the node for "scientist" would be updated based on its relationship to "discover" and "particle," which helps refine its meaning in this context.

As the system processes the sentence, it continuously updates the relationships between nodes and adjusts the weights of the edges based on the context. This allows the system to calculate meaning in real-time, with meaning emerging from the interconnectedness of the graph rather than being inferred from isolated words or predefined rules.

4. Applications of GNNs in Meaning Calculation

By using GNNs, AI systems can:

- Handle ambiguous language: Because meaning emerges from the relationships between nodes, GNNs can process ambiguous or multi-layered meanings. For instance, a word with multiple meanings can adjust its internal representation based on the surrounding nodes, selecting the most contextually appropriate meaning.
- Scale to complex knowledge systems: GNNs are highly scalable, making them ideal for structuring large-scale knowledge systems, such as knowledge graphs or semantic networks, where meaning needs to be calculated across many interconnected concepts.
- Integrate real-time updates: Since GNNs can handle dynamic updates to the network, they are well-suited for systems that need to continuously

learn from new data. A GNN-based language system could incorporate new words, phrases, or concepts as they appear in real-time, ensuring that the system's understanding remains relevant and accurate.

VIII. Applications and Implications

A. AI and Machine Learning

1. Reducing Dependence on Large Datasets

One of the most significant challenges in contemporary AI and machine learning is the sheer volume of data required for models to learn language effectively. Current large-scale language models, like GPT or BERT, are data-hungry, requiring billions of words and sentences to generate coherent outputs. This reliance on vast datasets creates issues related to:

- Resource intensiveness: Training these models consumes enormous computational resources, requiring large server farms, substantial energy, and financial costs.
- Bias and overfitting: Large datasets often contain biases that are hard to filter out, leading to models that may inadvertently reproduce or amplify those biases.
- Generalization limitations: While these models perform well on the data they have seen, they can struggle with generalizing beyond their training set, especially when faced with novel contexts or language use.

An a priori calculable language system would offer a radical shift. Since meaning would be calculated directly from linguistic structures rather than inferred from vast amounts of data, models based on this system would require far fewer data points to achieve a deep understanding of language. Such models would be:

- More efficient: By embedding semantic rules directly into the structure of the language, the system would no longer need to infer meaning from countless examples. Instead, it would calculate meaning as language is constructed, significantly reducing the need for large training corpora.

- Less biased: With meaning derived from the inherent structure of language rather than from the idiosyncrasies of data, an a priori calculable language system would be less prone to bias, as it wouldn't rely on biased datasets for learning.
- More adaptive: These systems could generalize more easily because they wouldn't be limited by the specific contexts in which words and phrases appear during training. Since the meaning is calculated dynamically, the system could apply its rules to any new context, handling novel language use more effectively.

For example, rather than learning that the word "light" can mean both "illumination" and "not heavy" based on patterns in a dataset, the system would calculate these meanings based on its internal rules and the context of usage. This could result in models that are more capable of interpreting ambiguous or polysemous terms, and better able to adapt to unseen words or new linguistic phenomena.

2. More Efficient Training Processes

In addition to reducing the volume of data required for training, an a priori calculable language system could also lead to more efficient training processes overall. Current deep learning models often require multiple stages of fine-tuning, gradient descent optimization, and iterative updates, all of which contribute to the resource-intensive nature of modern AI. By embedding meaning calculation directly into the model, much of this fine-tuning could be bypassed.

Training would focus on refining the semantic rules rather than learning from countless examples. Once the rules are established, they can be applied across a range of contexts without the need for continuous retraining or massive computational updates.

Additionally, these systems would allow for transfer learning in more flexible ways. Since the rules of meaning calculation would be domain-independent, a model trained in one domain (e.g., legal text) could be more easily adapted to another domain (e.g., medical research) without needing significant additional training data. This flexibility would revolutionize fields like machine translation,

sentiment analysis, and text summarization, where the ability to rapidly adapt to new domains is crucial.

B. Cognitive Science and Neuroscience

1. Understanding Human Cognition

The development of an a priori calculable language system would have profound implications for how we understand human cognition, especially with regard to language processing and meaning construction. Cognitive scientists have long debated how humans encode, process, and retrieve meaning. Is meaning built up through experience (as in empiricist models) or is it innately structured (as in rationalist models)?

An AI model capable of calculating meaning directly from linguistic structures would lend support to theories that propose innate cognitive structures for meaning. For example, the nativist theory of language, championed by thinkers like Noam Chomsky, suggests that humans are born with an innate "language faculty" that guides language acquisition. This model aligns with the idea of meaning being calculable a priori, as it suggests that much of our ability to understand language may be driven by pre-existing cognitive structures, rather than by pure learning through experience.

- Testing linguistic theories: Cognitive scientists could use these AI models to test hypotheses about human language processing. For instance, they could examine whether certain semantic rules embedded in the AI match how children acquire language, or how adults process ambiguous language in real-time.
- Understanding meaning representation: If meaning is calculable a priori, this suggests that the brain may also use a similar system for representing meaning. Neuroscientists could explore whether specific neural circuits correspond to semantic rules and relations encoded in the AI, helping to map the neural basis of language and meaning construction.

2. Brain-Like Models of Meaning Generation

Neuroscientists have long sought to understand how the brain processes meaning in real-time. Current models often rely on the idea that meaning emerges from neural networks that process sensory inputs and retrieve stored knowledge. However, if an a priori calculable system is developed, it could lead to a new understanding of how the brain generates meaning dynamically.

- Neural analogues: The rules of meaning calculation in this AI system might have direct analogues in the neural architecture of the brain. Just as the AI calculates meaning from linguistic structures, the brain may similarly calculate meaning from neuronal interactions and contextual cues, dynamically updating meaning as new sensory inputs arrive.
- Language and memory: One of the most complex aspects of human cognition is how language interacts with memory. An AI system that calculates meaning from structural rules might help model how the brain integrates short-term linguistic inputs with long-term semantic knowledge, offering insights into how we understand ongoing conversations, follow narratives, and recall information.

This could pave the way for brain-inspired AI models that not only simulate human cognition but also offer deeper insights into how meaning is constructed, represented, and manipulated in the brain. These models would potentially lead to advancements in neurological rehabilitation, education, and brain-machine interfaces by providing a clearer picture of how meaning emerges from cognitive processes.

C. Future Communication Systems

1. Efficient Communication with Fewer Symbols

One of the most exciting applications of an a priori calculable language system lies in the potential for designing new communication systems that operate with fewer symbols but convey greater semantic depth. Traditional communication relies on the transmission of long strings of symbols (words, letters, or sounds) to convey meaning. However, with a calculable language system, we could develop

systems where each symbol or combination of symbols is packed with more intrinsic meaning, reducing the amount of information that needs to be transmitted.

For instance:

- Condensed language: A sentence that requires ten words to convey in English could be encoded into just a few symbols in a calculable system, where each symbol inherently contains layers of meaning based on its structure and relationships. This could drastically reduce bandwidth requirements for communication, making it possible to convey complex ideas with minimal data transmission.
- Context-sensitive compression: In these systems, contextual information could be used to further compress the message. For example, in a known context (such as discussing a familiar topic), entire concepts could be encoded into a single symbol, with the receiver understanding the broader meaning based on shared knowledge of the subject matter.

This could revolutionize fields like:

- Human communication: People could communicate with greater efficiency, especially in environments where bandwidth is limited (e.g., space missions, remote regions, etc.). A single symbol could convey a paragraph's worth of information, allowing for more succinct and precise communication.
- Machine communication: Machines, especially in the context of Internet of Things (IoT) systems, could communicate more effectively with one another using these symbols, exchanging vast amounts of meaningful information with minimal data transmission.

2. Implications for Language Design

These new communication systems would also have implications for how we design languages in general. Traditional human languages have evolved over millennia to accommodate cultural, cognitive, and social needs. But with a calculable system, we could engineer languages from the ground up to be optimized for meaning transmission. Such a language could:

- Minimize ambiguity: By embedding meaning directly into linguistic structures, we could reduce the ambiguity often present in human languages, creating systems where every symbol is precise and contextually relevant.
- Facilitate human-AI interaction: As we move toward an era where humans and machines need to collaborate more closely, a calculable language system could serve as a universal interface, allowing humans and AI to communicate seamlessly. These systems could automatically translate between human languages and machine-optimized code, ensuring that communication remains efficient and mutually understandable.

3. Global Communication Networks

In an increasingly interconnected world, the need for efficient, universal communication systems is growing. A calculable language system could serve as the foundation for global communication networks, where individuals, machines, and systems from different cultural and linguistic backgrounds can communicate using a shared symbolic system. This could:

- Bridge language barriers: By offering a universal system of meaning calculation, we could overcome the limitations of current translation systems, allowing for instantaneous and seamless communication across different languages and dialects. Unlike current machine translation systems, which often struggle with nuances, idiomatic expressions, or cultural context, a calculable language system would rely on intrinsic meaning structures rather than statistical patterns. This would enable real-time translation that captures the full depth of meaning, ensuring clarity and understanding regardless of the languages involved.
- Enhance human-machine interaction: As the Internet of Things (IoT) and artificial intelligence continue to grow, machines from various domains and industries will increasingly need to communicate with one another, as well as with humans. A calculable language system could serve as a universal interface for machine-to-machine (M2M) and human-machine interaction, ensuring that communication is both efficient and contextually accurate. This would streamline operations in industries such as healthcare, logistics,

and smart cities, where quick and precise exchanges of information are critical.

- Facilitate global cooperation: In international diplomacy, science, and business, misunderstandings due to language barriers can lead to significant delays or conflicts. A universal, meaning-calculating language system would reduce these risks by enabling a shared linguistic framework for global collaboration. Such a system could be tailored to preserve cultural nuances while ensuring mutual understanding, fostering better cooperation in addressing global challenges like climate change, trade, and public health.
- In this vision of the future, a calculable language system would be at the core of global communication networks, allowing humans and machines to transcend linguistic and cultural differences, fostering a more connected and collaborative world.

IX. Challenges and Future Research

A. Feasibility of Implementation

While the idea of creating an a priori calculable language system presents revolutionary opportunities for artificial intelligence, cognitive science, and communication, there are significant challenges that must be addressed before such a system can become a reality. These challenges are both technical and theoretical in nature.

1. Technical Challenges

a. Complexity of Semantic Encoding:

- One of the key technical challenges is developing a system that can effectively encode semantic rules directly into linguistic structures. Meaning in natural language is inherently flexible and often context-dependent. Therefore, designing a system where meaning is calculated a priori requires encoding a vast array of semantic relationships that go beyond traditional rule-based systems or statistical language models.

- Current AI systems rely heavily on neural networks and deep learning architectures to process language by identifying statistical patterns in large datasets. While these models have demonstrated impressive capabilities, they lack the ability to understand or calculate meaning from first principles. Moving toward a system that can calculate meaning as language is constructed would likely require the development of new architectures that blend symbolic logic with adaptive neural networks—which is a significant technical hurdle.
- Moreover, building a system capable of handling real-time meaning calculation would involve integrating complex semantic relationships with syntactic rules in a way that is computationally efficient. This could lead to resource-intensiveness similar to or beyond current AI models, potentially requiring more advanced computational hardware or more refined algorithms to manage the complexity of meaning generation.

b. Context Sensitivity and Flexibility:

- One of the most difficult aspects of building such a system is addressing the issue of context sensitivity. In human language, meaning is highly dependent on context—words can have different meanings depending on their surrounding words, sentence structure, or even broader discourse. Current models struggle with maintaining consistent meaning across longer spans of text or in highly ambiguous scenarios.
- A system designed to calculate meaning from linguistic structures must be able to adapt dynamically to changing contexts. This would require the system to not only encode rules for meaning but also recognize when those rules should shift based on contextual factors. Developing an architecture that can efficiently handle this contextual dynamism is a major technical hurdle that would likely require breakthroughs in contextual modeling and dynamic updating mechanisms.

c. Scalability and Efficiency:

- Developing a system that calculates meaning a priori also raises questions of scalability. While the system may be able to handle small, controlled datasets or limited semantic domains, applying this model to more complex

and open-ended language use (as in real-world communication) would require it to scale effectively. How can a system that calculates meaning at the level of symbols, syntax, and context efficiently handle the immense variety of language used in human communication across different domains?

- Additionally, balancing efficiency and accuracy presents a challenge. Current systems already struggle with balancing computational load and language processing accuracy, and introducing real-time meaning calculation would increase the system's complexity. Developers will need to optimize algorithms to ensure that the system doesn’t become too computationally expensive or slow to operate in real-world applications.

d. Integration of Embodied Cognition:

- Embodied cognition, as discussed in earlier sections, links language meaning to physical experiences and sensory interactions with the world. Creating a language system that incorporates embodied meaning will require developing multimodal systems that integrate sensory input (visual, auditory, tactile, etc.) into the language processing model. While current AI systems have achieved some success in integrating multiple modalities, a fully integrated system where physical experiences directly influence meaning construction represents a significant technical challenge.

2. Theoretical Challenges

a. Formalizing Semantic Rules:

- A significant theoretical challenge lies in defining the semantic rules that govern the meaning of linguistic structures. While there are various theories of formal semantics (truth-conditional, model-theoretic, possible-worlds semantics), none of these frameworks currently allow for real-time meaning calculation that operates independently of external reference systems. Formalizing a set of semantic rules that are both comprehensive and adaptable is a complex task that requires deep linguistic and logical insights.
- Additionally, the theory of semantic entropy introduced earlier would need to be formalized and integrated into the system. This would require

defining clear metrics for measuring semantic complexity and richness, and these metrics must be computationally tractable while aligning with human intuitive understanding of meaning. Developing such a theory is a major theoretical challenge that would need to be resolved before implementation could proceed.

b. Accounting for Cognitive Variability:

- Human cognition is not uniform. People vary widely in how they understand and generate meaning, based on factors like cultural background, personal experience, education, and even neurological differences. One of the theoretical challenges in developing an a priori calculable language system is ensuring that the system is flexible enough to account for the variability in human cognition while still adhering to a core set of universal semantic principles.
- Cross-linguistic variability is also a challenge. Different languages have different syntactic and semantic structures, and a meaning-calculation system would need to be able to account for this diversity. How can a single system be designed to calculate meaning universally across all languages, or will such a system need to be language-specific?

B. Bridging Disparate Fields

The development of an a priori calculable language system would require collaboration across several fields, as the challenges span linguistics, cognitive science, information theory, computer science, and even philosophy. This level of interdisciplinary collaboration brings its own set of challenges, particularly in reconciling different methods, terminologies, and research goals. However, these collaborations are essential for developing a system that is both technically feasible and theoretically sound.

1. Collaboration Between Linguists and Computer Scientists

- Linguists bring expertise in understanding the rules and structures that govern natural languages, including syntax, semantics, pragmatics, and phonology. Their knowledge is crucial in formalizing the rules for semantic

grammar, defining the relationships between linguistic elements, and understanding how meaning is generated in different contexts. However, linguistic theory is often abstract and not immediately applicable to computational systems.

- Computer scientists and AI researchers, on the other hand, bring expertise in designing algorithms, building neural networks, and creating scalable, efficient systems. However, many AI researchers are not well-versed in the complexities of semantic theory or natural language processing beyond statistical methods.

Bridging this gap requires collaboration where both parties can work together to formalize linguistic rules in a way that is computationally feasible while ensuring that the nuances of meaning are preserved. For example, linguists might work with AI researchers to translate existing semantic theories into computable rules and structures, while computer scientists would focus on developing the algorithms and architectures to implement these rules efficiently.

2. Integrating Cognitive Science Insights into AI Systems

- Cognitive scientists bring valuable insights into how humans process and generate meaning in real-time. Cognitive models, such as predictive coding, can help inform the design of AI systems that mimic human meaning-making processes. However, cognitive models are often descriptive rather than prescriptive, meaning they describe what happens in the brain without specifying the exact mechanisms that could be used in a computational system.
- The challenge here is to take insights from cognitive psychology and neuroscience and apply them in a way that can inform the design of meaning-calculation algorithms. For example, cognitive scientists could work with AI researchers to build systems that simulate human-like processes for dynamically adjusting meaning based on context, ambiguity, and sensory input.
- Additionally, cognitive science insights could help resolve some of the theoretical challenges around embodied cognition, showing how sensory experiences are linked to linguistic meaning in human brains. This could

inform how to design AI systems that integrate multiple modalities to generate richer, more human-like meaning representations.

3. Contributions from Information Theory and Philosophy

- Information theorists are essential for developing a new theory of semantic entropy that accounts for the richness and depth of meaning, not just the unpredictability of symbols. These theorists can work on formulating metrics for semantic complexity, which would be critical for building systems that calculate meaning a priori. However, information theory has traditionally focused on symbol transmission rather than meaning generation, so expanding its scope to include semantics represents a novel challenge.
- Philosophers of language and epistemologists can contribute to the deeper questions surrounding meaning and knowledge representation. Issues like contextuality, ambiguity, and the relationship between syntax and semantics are areas where philosophy has made substantial contributions. Philosophers can help clarify the theoretical underpinnings of an a priori meaning-calculation system, ensuring that it aligns with broader theories of meaning and communication.

X. Conclusion

The exploration of an a priori calculable language of meaning represents a paradigm shift in how we approach language, communication, and artificial intelligence. The potential of such a system to reshape the future of AI, cognitive science, and communication is vast, touching on nearly every aspect of how meaning is constructed, interpreted, and transmitted across different contexts. By embedding meaning directly into the structural components of language—rather than relying on massive datasets or statistical inference—this approach offers the promise of systems that are more efficient, flexible, and adaptable than the language models we use today.

1. Reshaping the Future of AI

One of the most profound impacts of a calculable language system is in the realm of artificial intelligence. Current AI models are powerful but have significant limitations. They rely heavily on extensive datasets, can exhibit bias based on their training data, and often struggle with handling ambiguous or context-dependent meanings. A system that calculates meaning in real-time from the structure of language itself would offer:

- Reduced reliance on large datasets: By deriving meaning from inherent rules and relationships, AI systems would no longer require vast amounts of training data. This would lead to more efficient models, significantly lowering the computational and resource costs of AI training.
- Improved generalization: A system based on a priori meaning calculation would be better equipped to handle novel inputs, allowing AI models to generalize more effectively across domains, contexts, and languages.
- Less bias: Since meaning is calculated through internal structures rather than inferred from external data, the system would be less prone to reflecting the biases present in large datasets. This could lead to fairer and more ethically sound AI systems, better aligned with societal values.

In addition to these technical improvements, the development of AI systems that understand language through semantic grammar, dynamic ontologies, and graph neural networks could revolutionize human-machine interaction. AI systems would no longer be passive tools that process language; instead, they would become active agents capable of interpreting and engaging in complex conversations, adapting to new linguistic contexts in real-time, and providing nuanced responses that reflect a deeper understanding of meaning.

2. Advancing Cognitive Science and Neuroscience

The implications for cognitive science and neuroscience are equally profound. A system that can calculate meaning a priori would provide valuable insights into how the human brain processes and generates meaning. Such a system aligns with cognitive models like predictive coding, where meaning is constantly updated based on new inputs and the brain’s internal predictions. If AI can be

designed to mimic these processes, it may help cognitive scientists and neuroscientists:

- Test theories of language processing: The system could serve as a model for investigating how humans construct meaning in real-time. Researchers could test hypotheses about how semantic rules are encoded in the brain and how contextual shifts affect meaning.
- Explore brain-like models of meaning generation: By creating a system that simulates human-like processes of meaning construction, we could gain a deeper understanding of how the brain integrates language, sensory inputs, and memory to produce meaning. This could open up new research into areas like neural semantics and the neuroscience of language.
- Bridge AI and human cognition: The development of AI systems that generate meaning in a brain-like way could blur the line between artificial and human cognition, allowing for brain-computer interfaces that facilitate more natural, meaningful communication between humans and machines.

3. Revolutionizing Communication Systems

Perhaps one of the most exciting prospects of a calculable language system is its potential to revolutionize communication itself. By condensing meaning into fewer symbols, while simultaneously enriching the depth of each symbol's meaning, we could create:

- Highly efficient communication systems: These systems would be capable of transmitting complex ideas using minimal data, making communication faster and more efficient. This would be especially beneficial in environments where bandwidth is limited or where large-scale data transmission is impractical (e.g., deep-space missions, IoT networks).
- Universal communication interfaces: A calculable language system could serve as a universal language, capable of bridging linguistic and cultural barriers by representing meaning in a way that is accessible and understandable to all parties. This could lead to systems where humans and machines communicate seamlessly, using the same language for both machine learning and human conversation.

- Contextual meaning transmission: These systems would not only transmit raw data but also ensure that the contextual richness of meaning is preserved. This could drastically reduce misunderstandings and improve the precision of communication in both professional and personal domains.

4. Call for Interdisciplinary Research

Realizing the vision of an a priori calculable language system requires a massive interdisciplinary effort that integrates expertise from a wide range of fields. The challenges are vast—ranging from developing the underlying semantic theories and computational architectures to designing scalable and efficient systems capable of real-time meaning calculation. However, the potential rewards are equally vast.

- Linguistics: Linguists will need to formalize semantic rules that account for meaning in a calculable, structured way. These rules must be adaptable to different languages and contexts, requiring collaboration between experts in syntax, semantics, and pragmatics.
- Cognitive science and neuroscience: Researchers in these fields will play a critical role in providing insights into how meaning is constructed in the brain and how these processes can be mirrored in AI. Understanding the neural basis of meaning could inform the development of brain-like AI systems that simulate human cognitive processes.
- AI and computer science: AI researchers will be responsible for developing the computational architectures—such as neural-symbolic systems, graph neural networks, and dynamic ontologies—that can calculate meaning efficiently. Bridging the gap between symbolic logic and neural networks will be essential to building systems that can process meaning in real-time.
- Information theory: Information theorists will need to develop new metrics for semantic entropy and semantic complexity, providing a foundation for understanding how meaning can be measured and optimized in a language system.
- Philosophy: Philosophers of language and epistemologists will be essential for addressing deeper questions about meaning, context, and knowledge representation. Their work can help ensure that the system is grounded in

sound theoretical principles and that the implications for human communication are thoroughly understood.

5. A Unified Future of Meaning

In conclusion, the development of a language system that calculates meaning a priori offers a unified vision for the future of AI, communication, and human cognition. By embedding meaning directly into language structures, we can create systems that are more efficient, adaptable, and contextually aware. Such systems could transform the way we communicate, think, and interact with machines, creating a world where meaning is calculated with precision and depth, enhancing both human and machine capabilities.

The road ahead is challenging, but the potential rewards are immense. By fostering collaboration across disciplines, we can build the theoretical and practical tools necessary to realize this vision, opening the door to a new era of intelligent communication and cognitive modeling.

References

1. Shannon, C. E. (1948). "A Mathematical Theory of Communication." *Bell System Technical Journal*, 27(3), 379-423. DOI:10.1002/j.1538-7305.1948.tb01338.x
    - The foundational work on information theory, defining the concept of entropy and its application to communication systems.
2. Chomsky, N. (1957). *Syntactic Structures*. Mouton.
    - A seminal work in linguistics, introducing the concept of a generative grammar and arguing for an innate structure to language.
3. Lakoff, G., & Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
    - Introduces the theory of conceptual metaphors, linking abstract thought to bodily experiences and providing insights into embodied cognition.
4. Goller, C., Kuchler, A., & Sperl, D. (2020). "Graph Neural Networks for Natural Language Processing: A Survey." *arXiv preprint*. arXiv:2012.15491.
    - A comprehensive survey of the use of graph neural networks in natural language processing and how relational structures can be used to model language.
5. Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.
    - Explores the concept of mental representations and symbolic systems as the basis for cognition, providing insights into the nature of meaning generation in the human mind.
6. Pezzulo, G. (2008). "Coordinating with the Future: The Anticipatory Nature of Representation." *Minds and Machines*, 18(2), 179-225. DOI:10.1007/s11023-008-9095-5
    - Discusses the theory of predictive coding and its implications for anticipatory representation in cognitive systems.

7. Garcez, A., Gori, M., Lamb, L. C., Serafini, L., Spranger, M., & Tran, S. N. (2019). "Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning." *Artificial Intelligence*, 278, 139-160. DOI:10.1016/j.artint.2019.08.011
    - A detailed exploration of neural-symbolic systems, integrating the strengths of neural networks with symbolic logic for reasoning and meaning computation.
8. Piantadosi, S. T. (2014). "A Logical Language for Cognitive Science." *Behavioral and Brain Sciences*, 37(3), 355-356. DOI:10.1017/S0140525X13004011
    - Proposes a logical framework for understanding language and cognition, focusing on the relationship between symbols, logic, and meaning.
9. Baroni, M., & Lenci, A. (2010). "Distributional Memory: A General Framework for Corpus-based Semantics." *Computational Linguistics*, 36(4), 673-721. DOI:10.1162/coli_a_00016
    - Discusses the use of distributional semantics in NLP models, offering a contrast to symbolic approaches and highlighting the limitations of current models in capturing deep meaning.
10. Smolensky, P., & Legendre, G. (2006). *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*. MIT Press.
    - Explores the integration of neural computation with linguistic theory, providing insights into how deep meaning might be calculated in real-time through optimized grammar systems.
11. Marletto, C., & Deutsch, D. (2014). "Constructor Theory of Information." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2174), 20140540. DOI:10.1098/rspa.2014.0540
    - A theoretical framework for understanding information and computation in a physical context, useful for exploring how meaning could be encoded in a calculable system.

12. Goldstein, R. N. (2006). "Quantum Approaches to Semantic Information Processing." *Quantum Information Processing*, 5(4), 267-285. DOI:10.1007/s11128-006-0035-3
    - Discusses the potential of quantum information theory to model semantic processing, offering parallels between quantum superposition and layered meaning in language.
13. Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. MIT Press.
    - An examination of embodied cognition and how our physical interactions with the world shape our understanding of meaning, applicable to the design of language systems grounded in sensory experience.
14. Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press.
    - A comprehensive exploration of how meaning is generated through the interaction of grammar, cognition, and sensory experiences, with relevance to the design of calculable language systems.
15. Vinyals, O., & Le, Q. (2015). "A Neural Conversational Model." *arXiv preprint*. arXiv:1506.05869.
    - Presents early work on neural conversational models, providing a foundation for understanding how AI systems process meaning in dialogue, relevant to the future development of more advanced, meaning-calculating systems.